

InDetail Paper by Bloor

Author **Philip Howard and Daniel Howard**

Publish date **March 2021**

Software AG CONNX

“

**Software AG CONNX
needs to be considered
both in terms of its
capabilities as a generic
non-mainframe data
integration offering and
as one that supports
specific mainframe and
legacy functionality
either in conjunction with
Adabas or not.**

”

Executive summary

Software AG CONNX provides mainframe and legacy-oriented data integration, connectivity, data virtualisation and replication capabilities. It was originally developed by CONNX Solutions, a company that was acquired by Software AG in 2016 to support its “Adabas and Natural 2050+” agenda. It is complementary to the company’s webMethods EntireX offering, which offers application integration capabilities and webMethods ApplinX which offers screen-based integration. Between them, webMethods and CONNX enable two go-to-market strategies adopted by Software AG’s Adabas and Natural business unit, where the first of these is to maintain and transform the existing Adabas and Natural user base, and the second is to free mainframe users from the constraints of legacy environments, and to win new clients through data and mainframe integration that will enable mainframe environments to participate fully in modern hybrid and multi-cloud environments. In this paper we focus specifically on the capabilities of CONNX.

Key findings

In the opinion of Bloor Research, the following represent the key facts of which prospective users should be aware:

- CONNX offers three main sets of capabilities: data access (allowing applications to access supported data sources in real-time), data virtualisation (allowing federated query access across data sources, typically either from or including mainframe environments) and data movement (classical extract, transform and load [ETL], as well as extract, load and transform [ELT], and/or change data capture).
- CONNX provides native SQL access to more than 150 data sources. These are not limited to mainframe or legacy environments but extend to platforms such as Hadoop and databases such as VoltDB, MongoDB, AWS RedShift, Azure DB and Snowflake.

- Many of the CONNX mainframe and legacy connectors are unique and are re-sold by companies such as Informatica, Oracle, SAS and IRI.
- The CONNX DB Adapter automatically converts non-relational data so that it is accessible via SQL in real-time. This means that you can use any standard business intelligence or visualisation tool (Tableau, Qlik, Looker, Power BI, SAS and so on) in conjunction with CONNX. There is also an Excel plug-in available.
- CONNX ships with more than 600 pre-built transformations. When you are developing data pipelines you can either write SQL or there is a graphical user interface that supports a drag-and-drop paradigm and from which the software will generate SQL for you.
- Software AG has not yet implemented machine learning within CONNX to assist in building data integration pipelines (for example, suggesting next best transformation), though we understand that this is on the road map.
- CONNX has metadata management (which includes an understanding of COBOL copybooks) built-in, along with security at the user, group and application levels defined against database rows and/or columns.

The bottom line

The trend in the integration market is for vendors to offer all integration capabilities through a single platform. This includes not just database access but also Internet of Things (IoT) connectivity and integration, application integration, and API-based integration and management. In the case of Software AG these capabilities are spread across multiple products: CONNX, various webMethods offerings, and Cumulocity (for IoT connectivity). As we shall see, Software AG has a clear strategy with respect to how CONNX and webMethods

“
Software AG CONNX provides mainframe and legacy-oriented data integration, connectivity, data virtualisation and replication capabilities.
”



...a major differentiator for the company is that it has a depth of understanding of mainframe and legacy resources that few, if any, other companies can match...



work in conjunction, though this is less obvious with respect to Cumulocity. Leaving that aside, as a stand-alone data integration platform for accessing and leveraging a variety of data sources, CONNX has a lot to commend it. It is particularly noteworthy that Software AG has installed CONNX for new clients that are not mainframe users. Given that a major differentiator for the company is that it has a depth of understanding of mainframe and legacy resources that few, if any, other companies can match, it is heartening that even non-mainframe users are interested – and willing – to implement CONNX.

The context

Before discussing Software AG's CONNX specifically, it is worth considering the context within which it exists. There are two aspects to this: firstly, the broader Software AG strategy and, secondly, how CONNX fits into the company's plans for modernizing mainframe and other legacy environments. As far as the former is concerned, see [Figure 1](#), which provides a marketecture diagram illustrating the company's approach to enabling the connected enterprise. CONNX data integration plays a role in API management and hybrid integration and is a key element for Software AG's approach to mainframe evolution.

As far as mainframe/legacy environments are concerned, this has various aspects. IBM, for example, is promoting the idea of the mainframe as an analytics platform while a number of vendors are emphasising the use of modern development environments (DevOps) in conjunction with z/OS-based systems. Software AG also has a stake in this. At a high level it has adopted a "Adabas and Natural 2050+" agenda while more specifically, what it is aiming to offer to the wider market, is illustrated in [Figure 2](#). Here you can see that CONNX is just one part of this equation, working alongside webMethods ApplinX,

webMethods EntireX, webMethods API and Cumulocity (which provides connectivity to sensors and other Internet of Things-based devices and gateways).

What may not be clear from this diagram is that, under the covers, Software AG has adopted a microservices and (two-way) API-based approach in both the webMethods and CONNX products. The advantages of such methods are well-known and do not need to be rehearsed here, but Software AG is certainly ahead of some other data integration vendors in adopting these modern techniques, though we should also say that it lags behind some of those rivals when it comes to the implementation of machine learning to improve automation.




Figure 1: Marketecture of Software AG

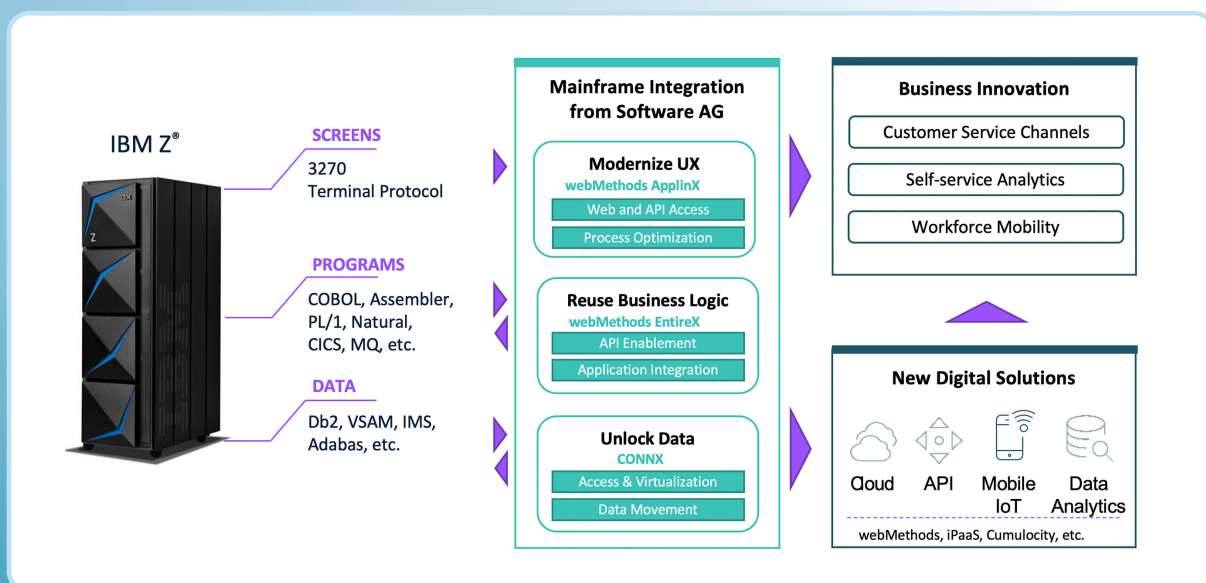


Figure 2: CONNX in context to the broader Software AG mainframe solution set

The product

Moving on to CONNX specifically, it serves a market beyond just mainframe to include legacy data sources on IBM i (AS/400), OpenVMS, legacy UNIX, and Windows platforms. Briefly, the way that it works is that you use the CONNX Client to develop your data pipelines (either via a visual drag-and-drop interface, for which there are more than 600 pre-built transformations

“adapters”, as opposed to “connectors”. In practice, many CONNX Adapters are simply connectors: you do not need anything special to access Microsoft SQL Server, for example. Conversely, if you want to access non-relational data sources then to do so via SQL requires some adaptation: converting SQL into a format suitable for accessing that source, which is what these adapters do. It is also worth noting that the SQL supported complies with the ANSI 92 standard. This is a lowest common denominator approach. While commonplace across data integration providers we would like to see support for more recent (less than three decades old!) versions of the ANSI standard, at least in specific cases. IBM, for example, supports ANSI 2016 across its database portfolio (Db2 on z/OS and LUW, Netezza and Informix) and users might like to leverage some of the more advanced features provided by that standard.

Data Access

Figure 4 represents the essentials of CONNX's Data Access use case: allowing third-party SQL access (from your desktop, the cloud or via mobile apps) to supported data sources, particularly on the mainframe, IBM i, OpenVMS and UNIX. You might also want the ability to access VSAM, IMS, Total, Adabas and a number of other legacy mainframe (and both IBM i and OpenVMS) environments. In this context it is worth commenting that the company's range of legacy adaptors is so extensive that they are resold by a number of otherwise competitive data integration vendors. However, the environment offered by CONNX is not limited to legacy/mainframe connectivity as fully two-thirds of its adapters are to non-legacy sources and targets.

Note that Data Access is two-way, supporting both read and write capability. This will be particularly useful if you want to write from Microsoft Excel to a mainframe source. For example, you could write from an Excel spreadsheet to a VSAM file with, we understand, just two clicks of your mouse.




Figure 3:
CONNX supports
Data Access, Data Virtualisation
and Data Movement




Figure 4:
Essentials of CONNX's
Data Access use case

delivered as part of the product, or by writing SQL code directly), while the Data Server and CORE components between them generate SQL (where necessary) and provide capabilities such as scheduling, alerting and so forth. Finally the DB Adapters (of which there are more than 150) provide SQL interfaces to relevant sources, including environments that do not natively support SQL. Security is applied at the user, group, and application level, and, within a relational environment at any rate, security applies at the row and/or column level. Metadata Management is provided to support data virtualisation while CONNX DataSync and InstantDbSync are used to support data movement use cases. And as shown in **Figure 3**, it is on the use cases – Data Access, Data Virtualisation and Data Movement – that CONNX supports, that we want to focus on, and not on the product's technical components.

However, before we do that, we need to discuss Software AG's use of the term

Data Virtualisation

While the Data Access use case is relatively simplistic, Data Virtualisation (shown in [Figure 5](#)) is more sophisticated. It allows you to access multiple data sources from within a single query. It works by allowing you to define a single SQL view across all relevant sources (even where those sources are non-relational). Your query is then distributed to those sources based on information held in the CONNX Data Dictionary, which holds relevant metadata about the sources being addressed, including such things as COBOL copybooks. Software AG talks about this allowing real-time querying and dashboarding and this will be true provided queries are not too large. Note that it is size that is specifically relevant here: complex queries (complexity is often conflated with size) shouldn't be a problem. From this perspective the key requirement is to minimise network traffic and the software does not include either a query optimiser or the sort of push-down optimisation that you would get from a pure-play data virtualisation tool. That said, many of the data sources supported by CONNX (Adabas, VSAM, IMS, RMS and so on) are not commonly, if ever, supported by pure-play providers, so the use of CONNX data virtualisation may be the only way that these data sources can participate within federated query processing. And, if you want to virtualise queries, some of which include mainframe/legacy sources and some of which do not, then you may be best served by leveraging a single solution such as the one that CONNX provides.

Note too, that CONNX only addresses databases and files: if you wanted to include data from, for example, an SAP ERP solution, then CONNX can get to the underlying data (historically, SAP ran on Oracle, Db2 and SQL Server), but users may want to combine both data integration and API-based approaches. Thus, you might use webMethods APIs to control what data is consumed based on business logic and security, with that data being passed to a staging area database for where CONNX picks up the data to perform ELT to a data lake, for example.

Data Movement

Data Movement represents, arguably, CONNX's primary use case. It is, after all, the bread and butter of data integration tools. As can be seen in [Figure 6](#), CONNX provides conventional extract, transform and load (ETL) capabilities though, not shown on this diagram, is the fact that it also supports ELT. Change Data Capture (CDC) is available, as illustrated. The difference between DataSync and InstantDbSync is partly that the latter, as its name implies, offers a real-time capability, but also that DataSync is generic whereas InstantDbSync, which is log-based, is limited to relational data sources (and, moreover, a limited subset of relational databases) plus Adabas.

Note that CONNX is SQL-based. This means that it does not require any intermediate processing of transformations, as products from companies such as Informatica require, which helps to reduce costs. Conversely, it means that SQL needs to be re-generated whenever a source or target schema changes. This could be an area where automation through machine learning might help to reduce this administrative burden, as well supporting functions such as recommending the next best transformation to use. Software AG has machine learning on its roadmap but has not yet committed to any particular use thereof.

Particularly noteworthy is that Software AG has integrated CONNX with Snowflake to enable migration, and ongoing support, for that platform. In this context is also worth commenting that Snowflake, in contrast to some of its direct competitors, is very limited in its support for data virtualisation, so leveraging CONNX for this purpose in conjunction with Snowflake would be worth considering in appropriate circumstances.




Figure 5:
Data Virtualisation in CONNX




Figure 6:
ETL capabilities in CONNX

The vendor



Software AG, founded in 1969, is headquartered in Darmstadt, Germany and has offices worldwide with more than 4,500 employees.



Software AG, founded in 1969, is headquartered in Darmstadt, Germany and has offices worldwide with more than 4,500 employees. Notable recent acquisitions have been Apama (2013), Zementis and CONNX (both 2016), Cumulocity (2017) and TrendMiner (2018). The company originally made its name with the Adabas database, and subsequently with the Natural 4GL, and Software AG now has two business units: its Digital Business (including Cloud and IoT), and Adabas and Natural (A&N). The company is listed on the Frankfurt Stock Exchange and, according to its latest figures (fourth quarter 2020), both divisions are growing strongly, with annual figures reflecting bookings growth of 21% for the Digital Business and 33% for A&N. For the full year, product revenue was €671.1 million, which was 1.6 percent lower than 2019 but reflecting an increasing shift from perpetual licences to subscription-based pricing.

Software AG has an extensive partner programme and for CONNX (see www.softwareag.com/connx), relevant partners include, but are not limited to IBM, Microfocus, Microsoft, Tableau, Teradata, Informatica, IRI and Qlik. The product's customer base includes Nissan, MetLife, Pfizer, Philips, HPE, Du Pont, Boston Children's Hospital, and many others.

Conclusion

Software AG CONNX needs to be considered both in terms of its capabilities as a generic non-mainframe data integration offering and as one that supports specific mainframe and legacy functionality either in conjunction with Adabas or not.

As a generic offering CONNX is competitive. The biggest trends in data integration are a) to provide more automation through the use of machine learning and b) to expand beyond simple connectivity to databases and files and to include both applications and IoT-based devices and gateways. As far as a) is concerned, Software AG is a little late to the party. Against that must be put the fact that it supports far more data sources than many competitive products. With respect to b) Software AG provides application integration via webMethods and sensor and device connectivity through Cumulocity. These both have significant strengths in their own right, but while the company seems to be on top of its marketing with respect to CONNX in conjunction with webMethods, this is less obvious when it comes to Cumulocity.

In mainframe environments, in which we include IBM i (AS/400) and OpenVMS as well z/OS-based systems, Software AG will often have significant advantages over competitors. The same may well be the case with respect to legacy UNIX and Windows environments. However, this will be less true if you are using Db2 or even IMS because these are well-served through rival offerings. For other database and file-based environments CONNX offers significant advantages though this is somewhat lessened by the fact that the company's adapters are resold by other players in this space. The data virtualisation functionality provided will be a significant advantage here as this is not a capability typically offered by other suppliers. Finally, if you are a legacy database (including, but not limited to, Adabas) user then the additional capabilities that CONNX offers, especially with respect to Snowflake integration, mean that CONNX should always be on your shortlist when it comes to data integration capabilities.



...if you are a legacy database user then the additional capabilities that CONNX offers, especially with respect to Snowflake integration, mean that CONNX should always be on your shortlist when it comes to data integration capabilities.



FURTHER INFORMATION

Further information about this subject is available from www.BloorResearch.com/update/



About the authors

PHILIP HOWARD

Research Director, Information Management

Philip started in the computer industry way back in 1973 and has variously worked as a systems analyst, programmer and salesperson, as well as in marketing and product management, for a variety of companies including GEC Marconi, GPT, Philips Data Systems, Raytheon and NCR.

After a quarter of a century of not being his own boss Philip set up his own company in 1992 and his first client was Bloor Research (then ButlerBloor), with Philip working for the company as an associate analyst. His relationship with Bloor Research has continued since that time and he is now Research Director, focused on Information Management.

Information management includes anything that refers to the management, movement, governance and storage of data, as well as access to and analysis of that data. It involves diverse technologies that include (but are not limited to)

databases and data warehousing, data integration, data quality, master data management, data governance, data migration, metadata management, and data preparation and analytics.

In addition to the numerous reports Philip has written on behalf of Bloor Research, Philip was previously editor of both ***Application Development News*** and ***Operating System News*** on behalf of Cambridge Market Intelligence (CMI). He has also contributed to various magazines and written a number of reports published by companies such as CMI and The Financial Times.

Philip speaks regularly at conferences and other events throughout Europe and North America.

Away from work, Philip's primary leisure activities are canal boats, skiing, playing Bridge (at which he is a Life Master), and dining out.



DANIEL HOWARD

Senior Analyst:

Information Management and DevOps

Daniel started in the IT industry relatively recently, in only 2014. Following the completion of his Masters in Mathematics at the University of Bath, he started working as a developer and tester at IPL (now part of Civica Group). His work there included all manner of software and web development and testing, usually in an Agile environment and usually to a high standard, including a stint working at an 'innovation lab' at Nationwide.

In the summer of 2016, Daniel's father, Philip Howard, approached him with a piece of work that he thought would be enriched by the development and testing experience that Daniel could bring to the table. Shortly

afterward, Daniel left IPL to work for Bloor Research as a researcher and the rest (so far, at least) is history.

Daniel primarily (although by no means exclusively) works alongside his father, providing technical expertise, insight and the 'on-the-ground' perspective of a (former) developer, in the form of both verbal explanation and written articles. His area of research is principally DevOps, where his previous experience can be put to the most use, but he is increasingly branching into related areas.

Outside of work, Daniel enjoys latin and ballroom dancing, skiing, cooking and playing the guitar.

Bloor overview

Technology is enabling rapid business evolution. The opportunities are immense but if you do not adapt then you will not survive. So in the age of Mutable business Evolution is Essential to your success.

We'll show you the future and help you deliver it.

Bloor brings fresh technological thinking to help you navigate complex business situations, converting challenges into new opportunities for real growth, profitability and impact.

We provide actionable strategic insight through our innovative independent technology research, advisory and consulting services. We assist companies throughout their transformation journeys to stay relevant, bringing fresh thinking to complex business situations and turning challenges into new opportunities for real growth and profitability.

For over 25 years, Bloor has assisted companies to intelligently evolve: by embracing technology to adjust their strategies and achieve the best possible outcomes. At Bloor, we will help you challenge assumptions to consistently improve and succeed.

Copyright and disclaimer

This document is copyright © 2021 Bloor. No part of this publication may be reproduced by any method whatsoever without the prior consent of Bloor Research.

Due to the nature of this material, numerous hardware and software products have been mentioned by name. In the majority, if not all, of the cases, these product names are claimed as trademarks by the companies that manufacture the products. It is not Bloor Research's intent to claim these names or trademarks as our own. Likewise, company logos, graphics or screen shots have been reproduced with the consent of the owner and are subject to that owner's copyright.

Whilst every care has been taken in the preparation of this document to ensure that the information is correct, the publishers cannot accept responsibility for any errors or omissions.



Bloor Research International Ltd
20-22 Wenlock Road
LONDON N1 7GU
United Kingdom

tel: **+44 (0)1494 291 992**
web: **www.Bloorresearch.com**
email: **info@Bloorresearch.com**